

Performance optimisation at UT

Victor Eijkhout

ICL, University of Tennessee, Knoxville

TOPS scidac

25/26 January 2002

Livermore CA

Topics of optimisation

- Matrix structure detection

Topics of optimisation

- Matrix structure detection
- A^2x kernel

Topics of optimisation

- Matrix structure detection
- A^2x kernel
- Multigrid smoothers

Matrix structure detection

Observations:

- Certain parallel preconditioners imply physical domain partitioning (Block Jacobi, but not multicolour ILU)

Matrix structure detection

Observations:

- Certain parallel preconditioners imply physical domain partitioning (Block Jacobi, but not multicolour ILU)
- 'Natural' domain partitioning often acknowledges partitioning of the physics.

Matrix structure detection

Observations:

- Certain parallel preconditioners imply physical domain partitioning (Block Jacobi, but not multicolour ILU)
 - 'Natural' domain partitioning often acknowledges partitioning of the physics.
- ⇒ Let partitioning for parallel processing acknowledge the same structure.

Structure test

Re-engineering of level sets:

Structure test

Re-engineering of level sets:

- first point connects to first point of next set

Structure test

Re-engineering of level sets:

- first point connects to first point of next set
- last point connects to last point of next set

Structure test


Re-engineering of level sets:

- first point connects to first point of next set
- last point connects to last point of next set
- first point does not connect to last point of another set


Structure test

Re-engineering of level sets:

- first point connects to first point of next set
- last point connects to last point of next set
- first point does not connect to last point of another set



Structure detection example



State of the work

Parallel implementation in Petsc

can be made more efficient by adding new Petsc primitives

State of the work


Parallel implementation in Petsc

can be made more efficient by adding new Petsc primitives

Small number of tests done

more to be done

comparison with Chaco &c.



A^2x kernel

(work for TSI scidac)

- Matrix is direct product of block diagonal and tridiagonal

A^2x kernel

(work for TSI scidac)

- Matrix is direct product of block diagonal and tridiagonal
- ADI preconditioner \Rightarrow solve many small dense systems

A^2x kernel

(work for TSI scidac)

- Matrix is direct product of block diagonal and tridiagonal
- ADI preconditioner \Rightarrow solve many small dense systems
- Solution of small (30–3k) dense systems by iterative method.
well-conditioned, so low number of iterations.

Iterative solution

- Formulate as left-preconditioned method

$$A = (D - E) \equiv D(I - N), \quad M^{-1} = (I + N)D^{-1}$$

$$\Rightarrow M^{-1}A = (I - N^2)$$

Iterative solution

- Formulate as left-preconditioned method

$$A = (D - E) \equiv D(I - N), \quad M^{-1} = (I + N)D^{-1}$$

$$\Rightarrow M^{-1}A = (I - N^2)$$

- depends on efficiency of $y = N^2x$
can we beat twice-gemv?

Efficiency of A^2x kernel

Why we can beat twice-gemv:

Efficiency of A^2x kernel

Why we can beat twice-gemv:

- Reuse of matrix

Efficiency of A^2x kernel

Why we can beat twice-gemv:

- Reuse of matrix
- Possible elimination of intermediate result

Efficiency of A^2x kernel

Why we can beat twice-gemv:

- Reuse of matrix
- Possible elimination of intermediate result
- Atlas gemv is optimised for out-of-cache; we possibly operate in-cache

Recursive approach to out-of-cache

case

Let

$$A = \begin{pmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{pmatrix}$$

then

$$y_1 = A_{11}(A_{11}x_1 + A_{12}x_2) + A_{12}(A_{21}x_1 + A_{22}x_2)$$

$$y_2 = A_{21}(A_{11}x_1 + A_{12}x_2) + A_{22}(A_{21}x_1 + A_{22}x_2)$$

Recursive approach continued

Introduce $t = Ax$ and localise application of A_{11} ,
 A_{22} :

$$\begin{Bmatrix} t_1 \\ t_2 \end{Bmatrix} = \begin{Bmatrix} A_{12}x_2 \\ A_{21}x_1 \end{Bmatrix}$$

$$y_1 = A_{11}\hat{t}_1, \quad \hat{t}_1 = A_{11}x_1 + t_1$$

$$y_2 = A_{22}\hat{t}_2, \quad \hat{t}_2 = A_{22}x_2 + t_2$$

$$\begin{Bmatrix} y_1 \\ y_2 \end{Bmatrix} + = \begin{Bmatrix} A_{12}\hat{t}_2 \\ A_{21}\hat{t}_1 \end{Bmatrix}$$

New kernel for recursive A^2x

Instructions involving reuse:

$$y_1 = A_{11}\hat{t}_1, \quad \hat{t}_1 = A_{11}x_1 + t_1$$

$$y_2 = A_{22}\hat{t}_2, \quad \hat{t}_2 = A_{22}x_2 + t_2$$

New kernel for recursive A^2x

Instructions involving reuse:

$$y_1 = A_{11}\hat{t}_1, \quad \hat{t}_1 = A_{11}x_1 + t_1$$

$$y_2 = A_{22}\hat{t}_2, \quad \hat{t}_2 = A_{22}x_2 + t_2$$

Double matrix vector product

$[y, s] = \text{m2v}(A, t, x) :$

$$s = Ax + t, \quad y = As$$

Example

On ev6:

	32	68	92	128
blas	356	415	395	389
unroll4	439	450	432	322

Example

On ev6:

	32	68	92	128
blas	356	415	395	389
unroll4	439	450	432	322
	64		96	
unroll4	386		399	

Multigrid smoothers

- Scalar optimisation (Atlas techniques); with Jun Ding.

Multigrid smoothers

- Scalar optimisation (Atlas techniques); with Jun Ding.
- Mathematical optimisation

construct CG iterates from GS iterates

does this pay?

other spectrum-adaptive method?

Summary

- Optimisation of dense and sparse kernels
- Optimisation: uni-processor and distributed
- Optimisation through Intelligent adaptation